Markus Furendal, Dept. of Political Science, Stockholm university and the Institute for Futures Studies

Collective ownership of AI

Abstract:
AI technology promises to be both the most socially important and the most profitable technology of a generation. At the same time, the control over – and profits from – the technology is highly concentrated to a handful of large tech companies. This chapter discusses whether bringing AI technology under collective ownership and control is an attractive way of counteracting this development.

It discusses *justice-based* rationales for collective ownership, such as the claim that, since the training of AI systems relies on a form of enclosing of the data commons, the value created by those systems should be fairly distributed. It also considers *democracy-based* rationales, like the suggestion that collective ownership is needed to ensure democratic control over the way this crucial technology is being developed and deployed. The paper also discusses possible forms of collective ownership, like publicly funded advanced AI research and democratically controlled AI companies. It concludes that the case for shifting to a model with collective ownership is the most compelling when based on the concern that the private model could come to concentrate economic and political power to such a degree that it threatens institutions designed to promote justice and democracy.

Keywords:
AI, AI governance, distributive justice, democracy, collective ownership

Artificial Intelligence (AI) is considered to be one of the most significant technological breakthroughs in several decades. Looking back at pivotal moments in the history of technology, like the shift from steam to electricity as the power source in industry, observers predict that AI technologies are set to boost productivity and create large profits. At the same time, there is a worry that AI-driven automation will displace a large share of the work force. In the medium to long run, this would not only reinforce the ongoing global trend of labour capturing a smaller share of total production and undermine the income-tax centred models for funding welfare state institutions, but also concentrate economic and political power to a few major AI developers. (Abbott and Bogenschneider 2018; Karabarbounis 2024; Acemoglu and Johnson 2023) Over and above these economic effects, there is an additional concern that AI systems may achieve abilities that would threaten the continued dominance – or even existence – of humans.

With stakes like these, it is crucial to ask who gets to control how AI is developed and deployed, and who gets to take part of the value it potentially creates. Currently, most powerful AI models, as well as the consumer-facing applications relying on those models, are controlled by a relatively small number of companies. Major decisions about what kind of AI to develop are made primarily by the boards of these companies, the managers and teams of engineers working in them and, in a formal sense, their shareholders. The revenues generated with AI technology go back into the companies, and their often wealthy founders and leaders. Call this the private ownership model.

Today, it is often assumed the private ownership model can be sufficiently fair and democratic, as long as there is sufficient regulation of – and ethical considerations within – the AI industry. The recently passed European AI act is designed to protect consumers and citizens from particularly risky applications of AI technology. Surprisingly many of Silicon Valley's business leaders claim to support efforts for redistributing some of the value the technology stands to create, including the idea of a Universal Basic Income, or even the unconditional access to computing power (Streitfield 2024). Still, as long as AI technology remains privately owned, citizens can only react to or try to

circumscribe significant decisions made in AI companies. Concerns like these reinvigorate long-standing issues around ownership and control over the economic sphere that go back at least to the century-old debate around markets and centralised planning (Adaman and Devine 1996). Given that we are still early in the trajectory of the commercial AI industry, and the economic effects have yet to materialize, we are currently at a unique point in time from which we can evaluate the benefits of different ways of organizing AI technology, and use governance to promote certain ideals.

This chapter hence explores the prospect of *collective ownership* of AI. The term is intended to capture different ways of changing the ownership structures of AI technology and companies, with the aim of increasing the degree of social control over the technology and the value it creates. The aim is, first, to sketch some of the ways that private ownership and control over AI could be complemented by different forms of collective alternatives, and second, to explore two sets of arguments for pursuing policies to promote forms of collective ownership, having to do with economic equality and justice, and political equality and democracy, respectively. The format of the chapter does not allow for a detailed discussion of economic theory or philosophical theories of property, and I can only briefly discuss some of the many objections that could be made to the positions sketched here. I do not claim that collective ownership would solve all of the problems associated with the private model, nor that it lacks problems of its own. Rather, the goal is simply to provide the first systematic discussion of the notion of collective ownership of AI, and prepare the ground for continued debate.

The upshot of the analysis is that, while there are several rationales for criticising the way the private ownership confers profits and power to the owners of AI technology, the solution does not necessarily involve abandoning the private model. The case for going further than mere regulation of the AI industry and encouraging some form of collective ownership seems more compelling, however, if we also consider that AI technology may come to concentrate economic and political power to such a degree that it threatens institutions designed to promote justice and democracy.

The next section (1) introduces the notions of private and collective ownership, discusses who controls AI technology today, and presents the analytical framework that guides the rest of the chapter. The following two sections (2, 3) discuss justice-based and democracy-based rationales with the aim of reconstructing the strongest version of each. Section 4 discusses some objections to the analysis I have presented, and a brief section 5 concludes.

1. Private and collective ownership of AI

In order to explore the notion of collectively owning AI we must consider, first, what distinguishes private and collective ownership models of productive assets, and second, in what ways it makes sense to think of AI as something that can be owned to begin with. In a basic sense, private ownership means that property rights are held by agents outside of the state. Conversely, collective ownership could be described as giving these rights to the state or some other publicly controlled collective entity. There is a common assumption that ownership matters because of the ways in which it confers control: In the traditional Marxist picture, for instance, the point of making the means of production collectively owned is to increase public control over the economy. Contrast this with a liberal ideal which rather seeks to maintain a distinction between the public and private spheres. Appealing to this ideal, liberals could oppose collective ownership because they want to prevent political power from exerting too much influence over what individuals do in the market.

In practice, however, it is difficult to describe productive assets as fully private or collective, since ownership rights are already circumscribed by regulation. For instance, tax laws give the state a right to some of the profits generated by property, and health and safety measures prevent factory owners from employing workers to use machines in particularly dangerous ways. Moreover, the relation between ownership and control is made more complex by the fact that ownership can be more or less dispersed. If the state is not democratically controlled, then 'collective' ownership of the means of production could be abused by a small political elite. Similarly, although the ownership over publicly traded business corporations is widely dispersed and, in a sense, collective, real control is arguably held not by the shareholders but rather the managers and executives who are employed by the firm and act on behalf of the owners. (Wright 2010, 114–16; Brouwer, Bennett, and Claassen 2022)

In the broadest sense, AI is clearly not the kind of property about which it makes sense to talk of ownership and the associated rights such as use, exclusion, and transfer. (Cf. Becker 1980) It is rather a set of mathematical and statistical methods or ideas, the fundamental character of which makes them unsuitable candidates even for intellectual property rights. For talk of owning or controlling AI to make sense, it should hence be considered as having to do with the concrete applications of these methods, or the constructed artefacts that rely on them. Although no one can own machine learning, for instance, patents could be awarded to an individual who develops a particular neural network-based system, or to the agency DARPA when it develops autonomous weapons systems using AI technology. Similarly, the Large Language Model GPT-4 could arguably be considered the property of OpenAI and under the current legal framework, they can license the use of the model in any way they see fit.[1]

Consider, next, the current ecosystem of AI. When it comes to the development of so-called foundation AI models and the consumer-facing applications that they power, there is an overwhelming dominance from the private sector. In 2023, 108 models were released by industry and only 28 by academia. Granted, a majority of foundation models are open-sourced and thus freely available in their entirety. Still, access to the most prominent models – including models from OpenAI, Google, Amazon, and Anthropic – is typically licensed and the underlying details of their design are proprietary information owned by the firms that developed them. (Maslej et al. 2024, 58)

This degree of dominance from industry has not always been the case, however. Before the breakthroughs in the machine learning approach in the mid 2010's, most notable AI models were developed in academic institutions (Maslej et al. 2024, 46) The reason industry has pulled ahead is generally attributed to the discovery that scaling up the models in terms of the amount of data they are trained on produces ground-breaking results, albeit in a highly resource-intensive process. This has made training cutting-edge models incredibly costly and necessitated massive capital investments. While Google's pioneering Transformer model released in 2017 is approximated to have cost around 900 dollars to train, the estimated training cost of their 2023 flagship model Gemini Ultra is around 190 million dollars. (Maslej et al. 2024, 64) It is not entirely clear just why developing AI technology is so expensive, but the standard explanation refers to the price of cloud computing services, which in turn are expensive in part because of bottlenecks in the production of the necessary hardware, including the chips called GPUs. (Clark 2023) One could speculate that there is also significant rent-seeking from the top AI talent, which are in short supply as more and

[1] The U.S. Patent Office recently rejected OpenAI's application to register the term "GPT" as a trademark, however, since it refers to a general type of technology ("generative pre-trained transformer") and not only to products developed by OpenAI. (David 2024)

more companies compete to hire AI engineers. Since the term "training costs" lacks a precise definition, however, and since companies keep most of the details around training runs secret, most estimates build on several levels of educated guesses.[2]

Even if we cannot neatly divide the economy into mutually exclusive categories of private and public ownership, it is nevertheless clear that ownership in the current AI ecosystem is best described as primarily private, since the right to decide how to develop and deploy AI – as well as the potential profits created by the use of the technology – are held primarily by non-state entities and individuals, largely excluding non-owners from direct control. AI development and deployment is primarily financed by private capital, and aims to provide private companies with AI models protected by existing property law.[3]

A few concrete examples of what a move toward collective forms of ownership could entail will help us when evaluating the rationales below. With regards to AI research, endeavours like the Manhattan project, the Apollo project and, more recently, Operation Warp Speed illustrate that it is not necessarily impossible for publicly funded research efforts to overcome the massive training costs described above, if there is political will and alignment with other policy goals. The 2023 executive order signed by the Biden administration in the U.S., for instance, created multiple National AI Research Institutes, and there is growing interest in Europe in creating a so-called 'CERN for AI'. (The White House 2023; Hoos and Irgens 2023). Policies around research funding and economic incentives can also steer private research in particular directions. The current patent-based system for innovation could, for instance, be complemented by a collectively funded and democratically controlled system of prizes for specific innovations. In such a system, collective decisions would stake out the preferred direction of AI development, AI companies would be able to recover their investments, but the resulting AI models would be freely available. (Erman and Furendal 2022, 286 f)

With regards to consumer-facing AI products and the value they create, we could rely on schemes like those above to create so-called public option complements to private AI companies' products, i.e. AI models that are "government-provided, quality-assured and universally available at a reasonable and fixed price". (Mazzucato and Li 2020, 1). Going a step further, strategies for moving toward collective ownership include setting up economic incentives for start-ups to be organised as worker cooperatives, or to shift the ownership of companies to employees through Employee Stock Ownership Plans and similar schemes. Another method involves collective capital institutions such as sovereign wealth funds or pension schemes buying stocks in publicly traded AI companies, or having companies offset some of their profits as shares in democratically controlled Citizens' Funds. A third and much more radical method is to simply nationalise AI-developing

[2] Another estimate for Gemini Ultra, which includes not only costs for the investment made in chips and the cost of running hardware operations but also the salaries of project members during the period of the training, puts the cost at 630 million dollars. This suggests that salaries make up more than two thirds of the total training costs. For a highly detailed description of these calculations, see https://colab.research.google.com/drive/1XEKlSo-3DCFp686yGOwwfS6_DEHsFimd?usp=sharing

[3] We should of course not overlook the importance of open-sourced models and AI tools developed in academia, but as we saw above, they currently make up a minor part of the AI ecosystem. Neither should we disregard the fact that, most, though not all, of the major AI-developing companies are publicly listed. On the other hand, some of them have a dual-class voting structure which ensure that the original founders retain a majority of the shareholder vote. The most extreme example of this is perhaps Meta, where the founder Mark Zuckerberg is the controlling shareholder, chair, as well as CEO. (Dignam 2020, 45)

companies, much like states have nationalised valuable industries and providers of public goods in previous eras.

Note that these are examples meant to illustrate the principled discussion below. It goes without saying that all such policies should be carefully evaluated, in terms of their benefits and problems, and their legitimacy and political feasibility, before they are possibly pursued. Still, the first step in this process involves analysing the philosophical and political values associated with each model. Hence, instead of focusing exclusively on particular proposals and their merits, this chapter will engage with the prior task of evaluating the rationales that speak in favour of shaping the ownership of AI in one way or the other through political decisions.

The rest of this chapter is devoted do discussing four rationales that have to do either with the value of justice, or the value of democracy. They should not be considered as four steps in an argument for collective ownership, but rather as four building blocks that can be modified and combined, and further developed into comprehensive arguments. The first rationale is primarily backward-looking, since it has to do with how AI technology is being developed. The other three are forward-looking, since they have to do with the effects AI technology can have on society.

2. Justice-based rationales

A commonly made critical observation in the age of generative AI is that, if it were not for the data that the models have been trained on, AI models would be nowhere near as powerful as they are. These data have, in turn, been generated through an unprecedented monitoring of individual behaviour through wearable hardware, software apps and social media. In addition, every last nook of the Internet has been surveyed by automatic crawlers, who harvest vast collections of text, video, audio, images, and music. (Zuboff 2019; Metz et al. 2024) A prominent justice-based argument for some form of collective ownership of AI departs from this observation, and says that changes to the private model are necessary to ensure the creators of these data are fairly compensated.

The first step in evaluating this argument is to distinguish between different forms of data production. The claim above seems relatively straightforward in cases where the works of particular artists have been used to train AI models without permission or compensation. When AI image generators became popular in 2022, for instance, users noticed that they often imitated the style of digital artist Greg Rutkowski, who for a long time had uploaded high-quality fantasy images, including extensive descriptions of each work, to the internet. This practice made his images particularly suited for scraping and collecting into databases, and gave his style an outsized influence over the way certain AI models generated images. (Heikkilä 2022) Even though AI companies have, in fact, treated everything that is publicly available on the internet as if it fell in the domain of "fair use" – and several of them are currently being sued by copyright holders who disagree (Brittain 2024) – it seems relatively easy to create a licensing scheme where the latter are compensated if their creations are in the training set of AI models. The company Adobe, for instance, recently solicited individuals to sell short annotated videos of people doing everyday things, such as walking or opening a door, to be used when training a text-to-video generator (Ford 2024). With a sufficiently robust licensing model in place, the problem of fair compensation for intentionally created data like these does not seem to require changes in the ownership of AI technology itself, although it would most likely change the economic viability of the companies developing these models.

It seems much more difficult, however, to design a system that compensates individuals for the use of what Shoshana Zuboff has called behavioural surplus data, that is, the by-products from living our lives in digital societies (Zuboff 2019, 69). Two fundamental problems prevent the licensing solution from being generalised to this case. The first is that data about any particular individual are, in fact, practically worthless. Unlike artists who can point to the influence their intellectual property has on particular AI models, the value of most of our individual data is negligible. Texts, images, and data about your behaviour become valuable only when they are aggregated into collections of data about large groups of individuals, or a sufficiently large corpus for AI models to train on.[4] This does not make a licensing scheme impossible, but it demonstrates that conceptualising the value of data as derivative of the labour expended in producing them, is likely to leave beneficiaries of such a scheme disappointed.

The second, and more fundamental, problem is that it seems difficult to set up the necessary contracts between individuals and AI developers that a licensing scheme requires. There are multiple reasons to question whether end user license agreements, for instance, constitute complete contracts in the technical sense in which economists use the term. (Wong, Duncan, and Lake 2024, 6) Most obviously, the data powering current AI-models are regulated by contracts signed long before either users or companies were familiar with the concept of training data. The internet forum Stack Overflow, for instance, has collected high-quality discussions around computer programming since 2008, and even profiled itself by banning the use of AI-generated text in the answers that users provide to each other's questions. This made it all the more surprising when the company recently struck a deal with major AI developers, selling access to their unique corpus of high-quality text produced by human users. When contributors complained, Stack Overflow pointed to a clause in its Terms of Service which declares content produced by users to be the property of the company. (Dave 2023; Grimm 2024) In response, users could reasonably insist that the agreement fails to live up to the clause about contractors having complete and perfect information given that many of them agreed to signing over rights to their answers without knowing that, in the future, this resource could be used in this way.

In light of these problems, I suggest that the most reasonable version of the rationale we are currently considering does not conceptualise data as individual property, but rather as a form of collective resource or 'digital commons' (Cf. Rosnay and Stalder 2020). A further benefit of this is that it allows us to recognise that the recent boom in AI capabilities is not just a result of training sets building on the work of living artists, but rather the fact that AI models consume and repeat all hitherto existing human culture, or at least that which is digitised and standardised in data sets. This is true beyond the domain of AI: Every current generation builds on the intellectual, cultural and material assets that previous generations have passed on. Not even the significant innovations in algorithmic processing which were necessary to unlock the value of the digital commons would have been possible without the common intellectual inheritance from previous generations. If data are digital commons, then the current boom in AI development is in fact a form of primitive accumulation, and the value AI creates is due to the fact that private companies are enclosing what was originally owned collectively.

Theorists have long argued that, when such enclosing happens, compensatory schemes are needed to ensure fairness. At the end of the 18th century, for instance, Thomas Paine suggested that, since the earth is a form of common property owned by all humans, the privatisation of natural resources creates an obligation to compensate those who no longer can access them. Specifically, he suggests

[4] There are companies founded on the idea of helping individuals to be compensated for the data they sell. Interestingly, however, most users can earn only between 5 and 15 $ per month doing so. (Wong 2023, 57)

this should take the form of a capital grant awarded to everyone when they turn 21. (Paine 2004; Cf. Vallentyne, Steiner, and Otsuka 2005) The pre-eminent proponent of Universal Basic Income in our age, Philippe van Parijs, has similarly argued that the distribution of the opportunity to engage in value-creating activities such as well-remunerated work or the invention of new technologies is neither natural nor necessarily fair. On his view, then, "…the taxes that fund a basic income are not levies on what was created out of nothing by today's producers, but rather fees to be paid by these producers for the privilege of using their personal benefit for what we have collectively received." (Van Parijs and Vanderborght 2017, 107) Proponents of data-owning democracy and digital socialism similarly argue that, because data is valuable and inherently collective in nature, it is not enough to simply regulate its use. The most promising way forward is rather to make citizens formal owners of data, in cooperatives, data trusts, or other institutions. (Muldoon 2022; Loi, Dehaye, and Hafen 2023; See also Chan, Bradley, and Rajkumar 2023).

In summary, the most convincing version of this backward-looking, justice-based rationale for collective ownership says that, since the data on which AI models are trained are a form of commons, and since the value created by enclosing the commons should be distributed fairly, AI-developing companies should share the benefits created by their technology more widely. Compared with the individual-level licensing scheme considered above, the collective approach is more inclusive when it comes to who should get compensated: one's right to compensation is not based on one's authorship of data that have gone into training sets, but rather on the collective ownership of the digital commons. On the other hand, the argument only demonstrates why fair compensation is necessary, and not that this requires changes to the private ownership model. As we just saw, it might be the case that redistributive schemes such as a universal capital grant or basic income are enough to satisfy this.

2.1. AI as a threat to justice-promoting institutions

A related but distinct justice-based rationale is forward-looking and holds that collective ownership is necessary since private ownership either threatens to undermine existing institutions that serve the interest of distributive justice, or prevents the creation of new institutions for fair compensation.

Consider, first, in what ways AI technology challenges *existing* institutions designed to promote a fair distribution. Today, many countries can be characterised as having more or less extensive welfare states, with some degree of taxation, redistribution, and provision of basic healthcare, education, public goods, and other resources, supported by the broadly liberal egalitarian view that individuals should be able to co-operate as equals under fair terms. One worry about AI technology is that automation threatens the viability of the welfare state not only by displacing workers – creating costs in terms of unemployment benefits – but also by undermining its funding model in a more fundamental way. Automation can indeed be defined as a shift from labour to capital as factors of production (Acemoglu and Restrepo 2018, 5 f), and neither unemployed workers nor the machines that replace them pay income tax. In the long run, economists suggest, automation may hence destabilise taxation schemes and macroeconomic models – which typically assume that the labour share of income is stable – and threaten a key source of government revenue needed for existing institutions. (Korinek 2020)

The second issue becomes apparent when we ask how we should raise the resources needed to fund *new* redistributive institutions, like a Universal Basic Income. One strategy would be to

conceptualise the willingness to give up parts of company revenue as a laudable Corporate Social Responsibility effort which would not only do good but might also translate into goodwill for the companies, as recently suggested by proponents of the so-called Windfall Clause. This proposal would have companies pledge that they will give away a share of their AI-boosted profits in the future, which could then be used to fund redistributive efforts (O'Keefe et al. 2020). Another strategy is to move away from taxing labour and ramp up taxes on different forms of capital, including the physical infrastructure on which AI technology depends, and the rents extracted from those who own the data on which AI models are trained. A related idea which has received some attention recently is to tax fixed resources, including land, much more heavily. (Korinek 2020)

How does collective ownership compare to these alternative solutions to AI-driven automation's impact on the welfare state? With regard to non-mandatory strategies like the former, I suspect most people are less optimistic about the intentions of key industry actors today compared with just a few years ago, in light of the backtracking of the more idealistic visions expressed by entities like OpenAI. A windfall clause thus seems bound to remain an empty promise. A change in the tax regime seems more promising, however, and could potentially be a politically feasible solution to the disruptive effects of automation.

There is nevertheless one important reason to consider forms of collective ownership as important complementary strategies, and it has to do with the normative justification of the redistributive scheme. Political interventions to redistribute resources often trigger strong intuitions that they are a form of expropriation. Many people feel that the income they perceive to be theirs should not be given to someone else unless they are presented with a convincing justification, or that the capital they have accumulated should be theirs to invest without the disincentive of taxation (Korinek 2020; cf. Murphy and Nagel 2002, who respond by questioning that all of our income is ours to begin with). The strategy of funding the welfare state through changes in ownership structures rather than taxation is a possible way to circumvent, as opposed to reject, this objection. For instance, instead of having the state tax AI companies' profits, it could create democratically controlled collective capital institutions like sovereign wealth funds, pension funds or a Citizens' Fund, in which companies are required to offset some of their profits as shares (Furendal and O'Neill 2024). A Universal Basic Income funded by such a fund would involve seeing citizens, not as parasites on the productive AI developers, but rather as collecting their dividends as legitimate owners of parts of the companies.

Now, which of these strategies is most feasible is clearly an empirical question which cannot be settled here. The point, however, is that we should not overlook the potential political force of the collective ownership strategy. A new land tax may be an attractive way of raising funds efficiently, but there is no logical connection between land value and the hardships suffered by displaced workers. A political story that begins with the digital commons-analysis reviewed above, however, and then describes those displaced by automation as rightful owners which should take part of the value created by the technology that displaced them, seems rather attractive. It makes sense to fund redistribution by tapping the productivity achieved through automation, and it is politically expedient to describe it as legitimate profit-sharing rather than remedial taxation.

3. Democracy-based rationales

Let us turn, finally, to discussing two democracy-based rationales. The first says that collective ownership is necessary for the governance of AI to be democratic. The second says that collective

ownership is a safeguard against private owners of AI technology undermining the preconditions for democratic governance more generally.

A key implication of the private ownership model is that most decisions concerning the future of AI are restricted to the small group of people working in tech companies, leaving non-owners without direct decision-making authority over the way AI is being built and used. Is this a problem? Voters can, after all, influence the shape of AI development indirectly – by electing politicians who make laws, supervise and set the boundaries for what these companies may do – as demonstrated by the recent U.S. executive order on AI and the EU AI act. Voters' interests are also represented if elected governments publicly fund AI research and steer the general direction of AI development. It might also be reassuring that many private companies have developed internal ethics guidelines regulating their work on AI, and that both employees and managers are concerned about the potential risks of the technology (Erman and Furendal 2022; Selling and Strimling 2023). Still, decision-makers in companies are not authorised by the general electorate, and have not been mandated to make significant decisions by the wider public who experience their consequences. On a reasonable understanding of democratic legitimacy, then, it is not sufficient that the interests of private AI decision-makers and the population align. A more robust process is needed to ensure key democratic values like representation and accountability. (Cf. Erman and Furendal 2024) Democratising AI, then, might very well require changes to the private model of ownership.[5]

This line of reasoning can find additional support from a classic argument for extending democratic principles from the political to the economic sphere, which appeals to the two being analogous. Robert Dahl, for instance, has argued that "if democracy is justified in governing the state, then it must also be justified in governing economic enterprises" (Dahl 1985, 111; Cf. Frega, Herzog, and Neuhäuser 2019) The basic democratic intuition that people should have a say about decisions that apply to them, could be interpreted to mean that indirect control over the economic sphere is not enough. The most radical interpretation rather says that, since ownership confers control, the full realization of this ideal requires the wider public become actual owners (Cf. Buller and Lawrence 2022). Although this is a general point about ownership of productive assets, it seems particularly important when it comes to AI development and deployment. In a well-functioning democracy, it could perhaps be realised by creating state-run AI enterprises with the mandate of voters to act in their interests. In a country with strong civil society actors or democratically run collective capital institutions, granting them formal influence by providing a seat on the boards of AI companies, could help to promote the interest of citizens. (Wright 2010, 116 f). None of these solutions are perfect, but they could potentially alleviate some of the democratic concerns around the private ownership model of today.

The second democracy-related argument centres not on how to increase the democratic control over AI-developing companies, but how the decisions of those companies could threaten existing democratic institutions. The last few years have seen extensive debate around the way AI can be used as a tool to more effectively interfere with elections and undermine trust and other key values of well-functioning democracies. (Cf. Coeckelbergh 2024) A distinct concern, however, is that the concentration of economic power in top AI-developing companies like Microsoft, Google and Meta make them politically influential actors in their own right, which limit the ability of citizens to make collective decisions about how to organise society. This is a specific instance of a more general problem in democracies, where high degrees of wealth concentration enable capital owners to extend their economic power into the political sphere. (Brouwer, Bennett, and Claassen 2022)

[5] Note that the term 'democratise AI' is notoriously vague in the debate, since it is often reduced to mean simply that access to AI technology ought to be wide and equal, i.e. 'democratic'. (Cf. Himmelreich 2023, 135; Lin 2024)

In some cases, it is relatively easy to recognize the exercise of such power, like when economically powerful AI companies lobby lawmakers to tailor legislation to suit their interests, which typically involves less strict regulation than what other civil society actors prefer, on the basis of promoting 'innovation'. (Perrigo 2023; Tallberg, Lundgren, and Geith 2024) Other times, however, less visible kinds of structural constraints created by economically powerful private actors make certain collective decisions so difficult or costly that they are pushed off the table. The global mobility of capital, for instance, creates a regulatory race to the bottom which restricts countries' ability to set the tax rates citizens prefer. Politicians in economically challenged areas often go to great lengths to accommodate requests for tax breaks and public subsidies from companies promising to create jobs. The economic potential of AI technology could very well reinforce the concentration of such structural power, granting whoever owns the technology significant sway over decisions that are formally made by democratic institutions. (Wright 2010, 83; Cohen 1989)

Importantly, unlike the problems we have discussed so far, this appears to be a distinct challenge which cannot be addressed through simple redistributive mechanisms but rather requires changes in ownership structures. That is, to counteract the kind of structural constraint on democratic decision-making that economically strong AI-developing companies constitute, it is not enough to simply redistribute the value created by AI technology. Neither is legislating around the kinds of AI they can develop, since there is a risk that all of the regulatory efforts discussed as alternatives to changes in ownership structures could be undermined if AI companies become too powerful. Hence, to institute forms of collective ownership of those companies would not only provide democratically selected representatives a seat at the table where the structurally consequential decisions are made. It could also help approximate the democratic ideal, by ensuring that a wider set of goals endorsed by the community would go into strategic decisions. (Cf. Furendal and O'Neill 2024, 316–20)

4 Objections and replies

There is undoubtedly a lot to react to in the discussion above, but two objections seem especially important to address here. The first rejects the account of property presented in the discussion on digital commons. The second argues that any move toward collective ownership is likely to come at a cost in terms of stifling innovation and economic growth.

First, a sceptic could point out that the analogy between the commonly owned earth and the digital commons fails because unlike natural resources, data are non-rival goods: The fact that an AI company scrapes the internet does not, in fact, prevent anyone else from doing the same. Moreover, on a standard Lockean understanding, the transformation of the digital commons to private property is justified by the fact that AI developers mix what they collect with their labour. (Cf. Nozick 1974) Raw data not only has to be cleaned and labelled, but significant resources are typically spent on having human workers finetune the way AI models treat the data, in a process known as reinforcement learning from human feedback. A better metaphor, the sceptic could conclude, would be that data are like a stream of water in which AI companies build hydroelectric plants and thereby gain property rights over the electricity produced, all the while allowing the water to flow freely for whoever else wants to build another power plant.

In response, we could nuance the initial claim to say that the enclosure of the digital commons is happening at a unique point in history, such that it will never again be as easy or cheap to extract the value contained in them. The main reason for this is that, as the Stack Overflow example

illustrates, there is currently an abundance of high-quality human-produced data online. Reports suggest, however, that leading AI companies have already incorporated most of this data and indeed reached the end of the data's potential, forcing them to scramble for additional resources. The main alternative is to train AI models on the output of other AI models, but research indicates that this may lead to 'model collapse', with quickly cascading errors. (Maslej et al. 2024, 52–55).[6] A second reason for believing that the last few years have seen a uniquely cheap and easy process of enclosure is that the ecosystem of the internet, in which most of this valuable data has been produced, is disrupted by the introduction of generative AI. First, there is a rapid increase in the absolute number of AI-generated websites, synthetic media and news services. If model collapse is a real threat, this part of the internet cannot be used to train new AI models. Second, as this undermines existing funding models for news sites and content creators, the relative share of AI-generated media increases even further. One reason why LLMs can produce excellent text is that they are trained on articles produced by journalists who are employed to write excellent texts. If generative AI disrupts the news industry, however, this kills the goose that lays the golden eggs. In this sense, even though the digital commons will always include high-quality data from the period up until around 2022, we cannot expect them to continue to be enriched by human contributions in the future.

The second objection worth discussing here says that, even though collective ownership might be attractive from the point of view of justice and democracy described above, any move away from the private ownership model will be costly in terms of other values: free and private enterprise could be said to be a requisite for scientific and economic progress, and given the value-creating potential of AI technology, it is not obvious that we should slow down progress. (Cf. Andreessen 2023)

A first response rejects the key assumption behind the objection, by emphasising the ways in which private research depends on public goods provided through collective effort and basic research in publicly funded academic settings. (Mazzucato 2014) In addition, even if private enterprise is currently pushing the research frontier, and academia struggles to attract top AI talent, we can easily imagine scenarios where the roles are the opposite. An international AI project of CERN or Apollo project size would arguably be able to compete both in money and prestige, not least because we could assume that many AI engineers are motivated not merely by receiving the highest possible salary but also by curiosity and the chance to push the limits of knowledge. It arguably matters less whether this happens in a private company or in a publicly funded and controlled project. (cf. Wright 2010, 63)

A second response would say that the objection relies on the implicit assumption that progress and development are good. We need not deny this, but could simply note that what the optimal pace of technological and economic progress is should arguably be up for democratic decisions. It might very well be the case that democratic decision-makers agree that AI progress should proceed at a rapid pace. A move toward collective ownership would ensure, however, that voices that are currently mostly overlooked also have a say in these deliberations. Some have raised concerns, for instance, about whether it is wise to devote resources and energy to developing AI models in light of other global challenges such as climate change. A more nuanced version of this concern says that it is a waste of resources to develop *multiple* cutting-edge AI models which are roughly as competent, as Silicon Valley companies currently do in fierce competition and at great economic and ecological costs. In a democratic deliberation, proponents of these views would have a say and

[6] Moreover, if the original data commons should be understood as collectively owned, then so should arguably the subsequent generations of AI-generated data produced by enclosing them

a chance to influence AI development, just as much as the progress-friendly voices which currently dominate.

5 Conclusion

This chapter has explored the somewhat radical proposal that the largely private ownership model that characterises contemporary AI development and deployment creates challenges for the ideals of distributive justice and democracy, and that this might lead us to consider moving toward forms of collective ownership of AI. There are undoubtedly more reasons that are worth considering, but which I have been forced to leave out.[7] A more extensive discussion would also weigh the goal of promoting democracy and justice against additional values, and address a number of well-rehearsed arguments in favour of the liberal model of maintaining a distinction between the public sector and the private economic sphere. The aim has nevertheless been to offer a first step in the careful evaluation needed to assess whether we should, indeed, move toward more collective ownership of AI.

The next step in such an evaluation also needs to engage in the difficult process of spelling out in more detail how collective ownership could work. Depending on which of the rationales discussed here we think is most convincing and important, different institutional efforts to create collective ownership of AI will have different benefits and drawbacks. Consider, for instance, that although worker cooperatives might promote the interests of employees in the AI industry, it would not help those who are displaced by automation. Relatedly, nationalising Silicon Valley's giants might perhaps empower Americans, but would still leave a vast majority of the world's population without any influence, even though the economic ramifications of AI automation are likely to be global. Great care and additional thought are thus needed when moving from the principled level on which I have focused to the more applied level of designing institutional solutions.

Acknowledgments:
[redacted for review]

References

Abbott, Ryan, and Bret Bogenschneider. 2018. "Should Robots Pay Taxes? Tax Policy in the Age of Automation." *Harv. L. &. Pol'y Rev.* 12:145–75.

Acemoglu, Daron, and Simon Johnson. 2023. *Power and Progress*. London: Basic Books.

Acemoglu, Daron, and Pascual Restrepo. 2018. "Artificial Intelligence, Automation and Work." Working Paper 24196, National Bureau of Economic Research, Cambridge, MA. http://www.nber.org/papers/w24196.

Adaman, Fikret, and Pat Devine. 1996. "The Economic Calculation Debate: Lessons for Socialists." *The Cambridge Journal of Economics* 20 (5): 523–37.

Altman, Sam Altman. 2021. "Moore's Law for Everything." March 16, 2021. https://moores.samaltman.com/.

Andreessen, Marc. 2023. "The Techno-Optimist Manifesto." October 16, 2023. https://a16z.com/the-techno-optimist-manifesto/.

[7] One worth mentioning has to do with the AI safety concern raised briefly in the introduction: if there is a risk that AI poses an existential risk, there are reasons to control access to the kind of compute needed to develop highly capable AI models.

Becker, Lawrence. 1980. "The Moral Basis of Property Rights." *NOMOS: American Society for Political and Legal Philosophy* 22:187–220.

Brittain, Blake. 2024. "Authors Sue Anthropic for Copyright Infringement over AI Training." *Reuters*, August 21, 2024, sec. Artificial Intelligence. https://www.reuters.com/technology/artificial-intelligence/authors-sue-anthropic-copyright-infringement-over-ai-training-2024-08-20/.

Brouwer, Huub, Michael Bennett, and Rutger Claassen. 2022. "Introduction: The Wealth Power Nexus." In *Wealth and Power*, 1st ed., 1–22. New York: Routledge.

Buller, Adrienne, and Mathew Lawrence. 2022. *Owning the Future*. London: Verso.

Chan, Alan, Herbie Bradley, and Nitarshan Rajkumar. 2023. "Reclaiming the Digital Commons: A Public Data Trust for Training Data." In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 855–68. Montréal, Canada: ACM.

Clark, Don. 2023. "How Nvidia Built a Competitive Moat Around A.I. Chips." *The New York Times*, August 21, 2023, https://www.nytimes.com/2023/08/21/technology/nvidia-ai-chips-gpu.html.

Coeckelbergh, Mark. 2024. *Why AI Undermines Democracy and What To Do About It*. Wiley.

Cohen, Joshua. 1989. "The Economic Basis of Deliberative Democracy." *Social Philosophy and Policy* 6 (2): 25–50.

Dahl, Robert A. 1985. *A Preface to Economic Democracy*. Quantum Books. Berkeley: University of California Press.

Dave, Paresh. 2023. "Stack Overflow Will Charge AI Giants for Training Data." *Wired*, April 20, 2023. https://www.wired.com/story/stack-overflow-will-charge-ai-giants-for-training-data/.

David, Emilia. 2024. "OpenAI Can't Register 'GPT' as a Trademark — Yet." The Verge. February 16, 2024. https://www.theverge.com/2024/2/16/24075304/trademark-pto-openai-gpt-deny.

Dignam, Alan. 2020. "Artificial Intelligence, Tech Corporate Governance and the Public Interest Regulatory Response." *Cambridge Journal of Regions, Economy and Society* 13 (1): 37–54.

Erman, Eva, and Markus Furendal. 2022. "The Global Governance of Artificial Intelligence: Some Normative Concerns." *Moral Philosophy and Politics* 9 (2): 267–91.

———. 2024. "Artificial Intelligence and the Political Legitimacy of Global Governance." *Political Studies* 72 (2): 421–41.

Ford, Brody. 2024. "Adobe Is Buying Videos for $3 Per Minute to Build AI Model." *Bloomberg.Com*, April 10, 2024. https://www.bloomberg.com/news/articles/2024-04-10/adobe-is-buying-video-clips-for-3-per-minute-to-build-ai-model.

Frega, Roberto, Lisa Herzog, and Christian Neuhäuser. 2019. "Workplace Democracy—The Recent Debate." *Philosophy Compass* 14 (4): e12574. https://doi.org/10.1111/phc3.12574.

Furendal, Markus, and Martin O'Neill. 2024. "Work, Justice, and Collective Capital Institutions: Revisiting Rudolf Meidner and the Case for Wage-Earner Funds." *Journal of Applied Philosophy* 41 (2): 306–29.

Gourevitch, Alex, and Lucas Stanczyk. 2018. "The Basic Income Illusion." *Catalyst* 1 (4).

Grimm, Dallin. 2024. "Stack Overflow Bans Users En Masse for Rebelling against OpenAI Partnership — Users Banned for Deleting Answers to Prevent Them Being Used to Train ChatGPT." *Tom's Hardware*, May 8, 2024. https://www.tomshardware.com/tech-industry/artificial-intelligence/stack-overflow-bans-users-en-masse-for-rebelling-against-openai-partnership-users-banned-for-deleting-answers-to-prevent-them-being-used-to-train-chatgpt.

Hagey, Keach, and Asa Fitch. 2024. "Sam Altman Seeks Trillions of Dollars to Reshape Business of Chips and AI." WSJ. February 9, 2024. https://www.wsj.com/tech/ai/sam-altman-seeks-trillions-of-dollars-to-reshape-business-of-chips-and-ai-89ab3db0.

Heikkilä, Melissa. 2022. "This Artist Is Dominating AI-Generated Art. And He's Not Happy about It." *MIT Technology Review*, September 16, 2022. https://www.technologyreview.com/2022/09/16/1059598/this-artist-is-dominating-ai-generated-art-and-hes-not-happy-about-it/.

Himmelreich, Johannes. 2023. "Against 'Democratizing AI.'" *AI & SOCIETY* 38 (4): 1333–46.

Hoos, Holger, and Morten Irgens. 2023. "CLAIRE Statement on Future of AI in Europe — 'AI Made in Europe' - Boost It or Lose It — 2023," June 26, 2023. https://claire-ai.org/wp-content/uploads/2023/07/CLAIRE-Statement-on-Future-of-AI-in-Europe-2023.pdf.

Karabarbounis, Loukas. 2024. "Perspectives on the Labor Share." *Journal of Economic Perspectives* 38 (2): 107–36.

Korinek, Anton. 2020. "Taxation and the Vanishing Labor Market in the Age of AI." *Ohio State Technology Law Journal* 16, no. 1: 244–57.

Lin, Ting-an. 2024. "'Democratizing AI' and the Concern of Algorithmic Injustice." *Philosophy & Technology* 37, article no. 103.

Loi, Michele, Paul-Olivier Dehaye, and Ernst Hafen. 2023. "Towards Rawlsian 'Property-Owning Democracy' through Personal Data Platform Cooperatives." *Critical Review of International Social and Political Philosophy* 26 (6): 769–87.

Maslej, Nestor, Loredana Fattorini, Raymond Perrault, Vanessa Parli, Anka Reuel, Erik Brynjolfsson, John Etchemendy, et al. 2024. "The AI Index 2024 Annual Report." Stanford: AI Index Steering Committee, Institute for Human-Centered AI, Stanford University. https://aiindex.stanford.edu/wp-content/uploads/2024/04/HAI_2024_AI-Index-Report.pdf.

Mazzucato, Mariana. 2014. *The Entrepreneurial State*. London: Anthem Press.

Mazzucato, Mariana, and Henry Lishi Li. 2020. "The Entrepreneurial State and Public Options: Socialising Risks and Rewards." Working Paper IIPP WP 2020—20. UCL Institute for Innovation and Public Purpose. https://www.ucl.ac.uk/bartlett/public-purpose/wp2020-20.

Metz, Cade, et al.. 2024. "How Tech Giants Cut Corners to Harvest Data for A.I." *The New York Times*, April 6, 2024, sec. Technology. https://www.nytimes.com/2024/04/06/technology/tech-giants-harvest-data-artificial-intelligence.html.

Muldoon, James. 2022. "Data-Owning Democracy or Digital Socialism?" *Critical Review of International Social and Political Philosophy*, September, 1–22.

Murphy, Liam B., and Thomas Nagel. 2002. *The Myth of Ownership*. Oxford: Oxford University Press.

Nozick, Robert. 1974. *Anarchy, State, and Utopia*. New York: Basic Books.

O'Keefe, Cullen, et al. 2020. "The Windfall Clause: Distributing the Benefits of AI." Centre for the Governance of AI Research Report. Oxford: Future of Humanity Institute. https://www.fhi.ox.ac.uk/windfallclause/.

Paine, Thomas. 2004. "Agrarian Justice (1797)." In *The Origins of Universal Grants: An Anthology of Historical Writings on Basic Capital and Basic Income*, edited by John Cunliffe and Guido Erreygers, 3–16. London: Palgrave Macmillan UK.

Perrigo, Billy. 2023. "OpenAI Lobbied E.U. to Water Down AI Regulation." Time. June 20, 2023. https://time.com/6288245/openai-eu-lobbying-ai-act/.

Rosnay, Mélanie Dulong de, and Felix Stalder. 2020. "Digital Commons." *Internet Policy Review* 9 (4).

Selling, Niels, and Pontus Strimling. 2023. "Liberal and Anti-Establishment: An Exploration of the Political Ideologies of American Tech Workers." *The Sociological Review* 71 (6): 1467–97.

Streitfeld, David. 2024 "Is It Silicon Valley's Job to Make Guaranteed Income a Reality?," *The New York Times*, July 16, 2024, sec. Technology, https://www.nytimes.com/2024/07/16/technology/ubi-openai-silicon-valley.html.

Tallberg, Jonas, Magnus Lundgren, and Johannes Geith. 2024. "AI Regulation in the European Union: Examining Non-State Actor Preferences." *Business and Politics*, February, 1–22.

The White House. 2023. "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence." https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/.

Vallentyne, Peter, Hillel Steiner, and Michael Otsuka. 2005. "Why Left-Libertarianism Is Not Incoherent, Indeterminate, or Irrelevant: A Reply to Fried." *Philosophy & Public Affairs* 33 (2): 201–15.

Van Parijs, Philippe, and Yannick Vanderborght. 2017. *Basic Income*. Cambridge: Harvard University Press.

Wong, Wendy H. 2023. *We, the Data*. Cambridge, Massachusetts: MIT Press.
Wong, Wendy H., Jamie Duncan, and David A. Lake. 2024. "Why Data about People Are so Hard to Govern." *Regulation & Governance*, April, rego.12591.
Wright, Erik Olin. 2010. *Envisioning Real Utopias*. London ; New York: Verso.
Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for the Human Future at the New Frontier of Power*. London: Profile Books.